



[Subscribe](#) (Full Service) [Register](#) (Limited Service, Free) [Login](#)

Search: ☐ The ACM Digital Library ☐ The Guide

SEARCH



[Feedback](#) [Report a problem](#) [Satisfaction survey](#)

High-speed policy-based packet forwarding using efficient multi-dimensional range matching

Full text Pdf (1.82 MB)

Source [Applications, Technologies, Architectures, and Protocols for Computer Communication archive](#)

Proceedings of the ACM SIGCOMM '98 conference on Applications, technologies, architectures, and protocols for computer communication [table of contents](#)

Vancouver, British Columbia, Canada

Pages: 203 - 214

Year of Publication: 1998

ISSN:0146-4833

[Also published in ...](#)

Authors [T. V. Lakshman](#)
[D. Stiliadis](#)

Sponsor [SIGCOMM](#): ACM Special Interest Group on Data Communication

Publisher ACM Press New York, NY, USA

Additional Information: [abstract](#) [references](#) [citations](#) [index terms](#) [collaborative colleagues](#) [peer to peer](#)

Tools and Actions: [Discussions](#) [Find similar Articles](#) [Review this Article](#)
[Save this Article to a Binder](#) [Display in BibTex Format](#)

DOI Bookmark: Use this link to bookmark this Article: <http://doi.acm.org/10.1145/285237.285283>
[What is a DOI?](#)

↑ ABSTRACT

The ability to provide differentiated services to users with widely varying requirements is becoming increasingly important, and Internet Service Providers would like to provide these differentiated services using the same shared network infrastructure. The key mechanism, that enables differentiation in a connectionless network, is the packet classification function that parses the headers of the packets, and after determining their context, classifies them based on administrative policies or real-time reservation decisions. Packet classification, however, is a complex operation that can become the bottleneck in routers that try to support gigabit link capacities. Hence, many proposals for differentiated services only require classification at lower speed edge routers and also avoid classification based on multiple fields in the packet header even if it might be advantageous to service providers. In this paper, we present new packet classification schemes that, with a worst-case and traffic-independent performance metric, can classify packets, by checking amongst a few thousand filtering rules, at rates of a million packets per second using range matches on more than 4 packet header fields. For a special case of classification in two dimensions, we present an algorithm that can handle more than 128K rules at these speeds in a traffic independent manner. We emphasize worst-case performance over average case performance because providing differentiated services requires intelligent queueing and scheduling of packets that precludes any significant queueing before the differentiating step (i.e., before packet classification). The presented filtering or classification schemes can be used to classify packets for security policy enforcement, applying resource management decisions, flow identification for RSVP reservations, multicast look-ups, and for source-destination and policy based routing. The scalability and performance of the algorithms have

been demonstrated by implementation and testing in a prototype system.

[Subscribe \(Full Service\)](#) [Register \(Limited Service, Free\)](#) [Login](#)Search: ☒ The ACM Digital Library ☐ The Guide **SEARCH**[Feedback](#) [Report a problem](#) [Satisfaction survey](#)

Packet classification on multiple fields

Full text Pdf (1.46 MB)

Source [Applications, Technologies, Architectures, and Protocols for Computer Communication archive](#)
Proceedings of the conference on Applications, technologies, architectures, and protocols for computer communication [table of contents](#)
Cambridge, Massachusetts, United States
Pages: 147 - 160
Year of Publication: 1999
ISBN: 1-58113-135-6
[Also published in ...](#)

Authors [Pankaj Gupta](#)
[Nick McKeown](#)

Sponsor [SIGCOMM: ACM Special Interest Group on Data Communication](#)

Publisher ACM Press New York, NY, USA

Additional Information: [abstract](#) [references](#) [citations](#) [index terms](#) [collaborative colleagues](#) [peer to peer](#)

Tools and Actions: [Discussions](#) [Find similar Articles](#) [Review this Article](#)
[Save this Article to a Binder](#) [Display in BibTex Format](#)

DOI Bookmark: Use this link to bookmark this Article: <http://doi.acm.org/10.1145/316188.316217>
[What is a DOI?](#)

↑ ABSTRACT

Routers classify packets to determine which flow they belong to, and to decide what service they should receive. Classification may, in general, be based on an arbitrary number of fields in the packet header. Performing classification quickly on an arbitrary number of fields is known to be difficult, and has poor worst-case performance. In this paper, we consider a number of classifiers taken from real networks. We find that the classifiers contain considerable structure and redundancy that can be exploited by the classification algorithm. In particular, we find that a simple multi-stage classification algorithm, called RFC (recursive flow classification), can classify 30 million packets per second in pipelined hardware, or one million packets per second in software.

↑ REFERENCES

Note: OCR errors may be found in this Reference List extracted from the full text article. ACM has opted to expose the complete List rather than only correct and linked references.

1 [Mikael Degermark , Andrej Brodnik , Svante Carlsson , Stephen Pink, Small forwarding tables for fast routing lookups, Proceedings of the ACM SIGCOMM '97 conference on Applications, technologies, architectures, and protocols for computer communication, p.3-14, September 14-18, 1997, Cannes, France](#)